arXiv:1703.06471v1 [cs.AI] 19 Mar 2017

Multi-Timescale, Gradient Descent, Temporal Difference Learning with Linear Options

Peeyush Kumar
agaron@uw.edu
University of Washington Seattle

Doina Precup
dprecup@cs.mcgill.ca
McGill University, Montreal

March 21, 2017

Abstract

Deliberating on large or continuous state spaces have been long standing challenges in reinforcement learning. Temporal Abstraction have somewhat made this possible, but efficiently planing using temporal abstraction still remains an issue. Moreover using spatial abstractions to learn policies for various situations at once while using temporal abstraction models is an open problem. We propose here an efficient algorithm which is convergent under linear function approximation while planning using temporally abstract actions. We show how this algorithm can be used along with randomly generated option models over multiple time scales to plan agents which need to act real time. Using these randomly generated option models over multiple time scales are shown to reduce number of decision epochs required to solve the given task, hence effectively reducing the time needed for deliberation.

1 Introduction

Decision making involves choosing among different actions over a multiple range of time scales. The use of higher level temporally extended actions allow an AI agent to solve an entire task in smaller number of decision epochs. Like many other AI systems, Temporal abstraction is often used in reinforcement learning to solve large problems efficiently. Various methods have been proposed to reduce computational burden, and reducing data complexity in learning, by allowing planning in AI systems to compute a good policy for many situations at once. In environments with continuously valued observations or others with a large number of states, some form of generalization is necessary. Linear methods are perhaps the most common and best understood class of generalization mechanisms. Linear methods encompass a wide spectrum of approaches including look-up tables, state aggregation methods, radial basis functions with fixed bases, etc. Frameworks for temporal abstractions do not usually contain an intrinsic structure to describe such generalizations. To achieve generalizations usually temporal abstractions are defined over parametric representation of the environment states, among which linear representations are the most dominant approach.

Since temporal effects are linked to the agents decision method, temporal abstraction framework is tightly coupled with controls. Markov Decision Process (MDP) models provide a formal method of planning in controlled dynamical systems (Sutton and Barto (1998)). The MDP models for planning which are based on primitive actions can be easily extended to SMDP (Semi Markov Decision Process) planning models which include temporally extended actions. The SMDP theory with models based on *options* framework, provides a formal method for representing temporally abstract information. An option $\mathcal{O}$ is a tuple $\{\mathcal{I}, \mu, \beta\}$, where $\mathcal{I}$ is the initiation set of the option, μ is the behavior policy of the option $\mathcal{O}$, β is the termination condition which defines the time scale of the option. The option models (μ and β) are either constructed by learning maximum-likelihood primitive models and computing exactly option models by composing primitive models or learning option models using intra-option learning, importance sampling, or eligibility traces.

Although the option models have been proved to be quite successful in continuous state spaces (one of the central advantages of temporal abstraction), the ability to plan using temporal abstraction methods is still a challenge when the number of states is large. This is in large partly because learning option models for large state spaces is very expensive. Moreover, if the system needs to act in real time, there is not much time for deliberation. In such cases the system needs to make use of many option models over multiple time scales. We propose here that instead of learning option models from the data, the system should rather make use of many option models over multiple time scales even if they are a lot in number. We use the linear option models (Sorg and Singh (2010)) for learning and planning. Learning a policy over many option models does not effect the time complexity because the linear methods scale linearly with the complexity of the value function. Linear option model extends

the option framework from a tabular representation to a linear representation. These models provide a general mechanism for providing temporal abstraction over parametric representation of the state space. Behavior of such a model is proved to be sensible while learning and planning. Although linear options provide great benefits towards time efficiency, the SMDP-LSTD algorithm used for policy evaluation using linear options has a quadratic complexity. This makes their convergence slow as compared to the conventional linear complexity of the TD algorithms.

Recently various temporal difference algorithms were proposed which are compatible with both linear function approximation and off-policy training using primitive action in MDP setting (GTD algorithm Sutton et al. (2001), GTD2 and TDC algorithms Sutton et al. (2009)). These algorithms were shown to converge to the TD fixed point with probability one. We propose a direct extension of the TDC algorithm to the SMDP setup. We show the convergence conditions for the SMDP-TDC algorithm using results from stochastic approximation theory.

2 Background

An MDP is a tuple (S, A, P, R, γ) consisting of a state set S, action set A, transition function $P : S \times A \times S \mapsto [0, 1]$, an expected reward function $R : S \times A \mapsto \mathbb{R}$, and a discount factor $\gamma \in [0, 1)$. Every MDP time step t, the agent takes an action a_t in state s_t, and the environment responds with a reward r_t and resulting next state $s'_t = s_{t+1}$. A policy $\pi : S \times \mapsto [0, 1]$ is a mapping from states to actions. An action-value function, $Q(s, a)$ is the expected discounted return obtained by taking action a in state s and following π thereafter. The Bellman equation defining the optimal action value function Q^* is:

$$Q^*(s, a) = R(s, a) + \gamma \sum_{s' \in S} P_a^{ss'} \max_{b \in A} Q^*(s', a)$$

A policy that is greedy with respect to a given action value function maps each state to $\arg\max_a Q(s, a)$. The value function $V(s)$ corresponding to a given action value function is the value of the greedy action from each state: $V^\pi(s) = \max_a Q^\pi(s, a)$. A policy that is greedy with respect to the optimal value function is the optimal policy.

2.1 Off Policy Temporal Difference Learning

Off Policy methods are an important classes of algorithms studied in Reinforcement Learning. Any controlled dynamical system is said to be off policy, if the behavior policy (the one which is being followed) is different from the policy being evaluated. Among various methods of learning in RL, temporal difference methods are undoubtedly the most important and novel ones. Temporal Difference methods have the advantage that they can learn directly from the raw data

without any requirement of the environment's model as well as they can bootstrap by updating estimates based in part on other learned estimates. Among the temporal difference (TD) methods, off policy TD control algorithm known as Q-learning is the most central one. Its simplest form, one-step Q-learning, is defined by

$$Q(s_t, a_t) = Q(s_t, a_t) + \alpha \left[r_{t+1} + \gamma \max_a Q(s_{t+1}, a) - Q(s_t, a_t) \right]$$

2.2 Linear Methods for MDPs

In reinforcement learning, linear methods have long been used for value function approximation. Instead of using tabular representation of states, states are represented in terms of n dimensional feature vectors $\phi(s) \in \mathbb{R}^n$. Linear methods approximate the value function for a policy π as linear combination of state feature vectors. $V^\pi(s) \approx \phi(s)^T \theta^\pi$, where θ^π is the vector parameter to be learned, $\theta \in \mathbb{R}^n$.

Unfortunately Q-learning does not show a good behavior when used with function approximations that are linear in the learned parameters. As a result convergence cannot be guaranteed for such methods (Baird (1995)). Several non-gradient-descent approaches to this problem have been developed, but none has been completely satisfactory . Second-order methods, such as LSTD (Bradtke et al. (1996), Boyan (2002)), can be guaranteed stable under general conditions, but their computational complexity is quadratic in the size of the state space. In recent years several gradient descent methods have been proposed (GTD, GTD2, and TDC algorithms). These methods are shown to have a linear computation complexity and linear memory requirements. Among these methods TDC shows the fastest rate of convergence. The update estimates for TDC are:

$$\theta_{k+1} = \theta_k + \alpha_k(\delta_k \phi_k - \gamma \phi'_k)(\phi_k^T w_k) \tag{1}$$

where α_k and β_k are the learning rate parameters, and w_k is updated as,

$$w_{k+1} = w_k + \beta_k(\delta_k - \phi_k^T w_k)\phi_k \tag{2}$$

and δ_k is the TD error

$$\delta_k = r_k + \gamma \theta_k^T \phi'_k - \theta_k^T \phi_k \tag{3}$$

It is proved that TDC algorithm converges to the TD fixed point with probability one.

2.3 Options and SMDP

Options are temporally extended actions. They can be thought of as fixed policies that can be invoked in certain set of states and which terminate according to a termination condition. Formally an option is a triple $(\mathcal{I}, \mu, \beta)$, where the

set $\mathcal{I} \subset \mathcal{S}$ is the initiation set of the option, $\mu : S \mapsto A$ is the option policy and $\beta : S \mapsto [0, 1]$ maps each state to the set $[0, 1]$ which is the probability of termination in each state $s \in S$. MDP along with the options describe a Semi Markov Decision Process (SMDP). SMDP is like a MDP, except the options take varying amounts of time. In each SMDP time step, the agent selects an option and follows the option's policy μ until termination. In any state, one has to choose an option from the set of options available in that state. When the system enters a new subset of the state space, a new set of options becomes available.

The Bellman equation for SMDP setup is a direct extension of the Bellman equation for the MDP case

$$Q^*_{\mathcal{O}}(s, o) = r^o_s + \sum_{s'} P^{ss'}_o \max_{o' \in \mathcal{O}'_s} Q^*_{\mathcal{O}}(s', o')$$

where r^o_s is the expected discounted reward obtained during the execution of the option

$$r^o_s = E\left[r_t + \cdots + \gamma^{k-1} r_{t+k-1} | s_t = s, o_t = o\right]$$

and $P^{ss'}_o$ is the option's transition model

Similarly the SMDP version of one step Q learning using options can be written as follows

$$Q(s, o) = Q(s, o) + \alpha \left[r + \gamma^k \max_{o' \in \mathcal{O}_{s'}} Q(s', o') - Q(s, o)\right]$$

where k is the number of time steps elapsed between s and s', r denotes the cumulative discounted reward until the termination of the option.

2.4 Linear Options

Linear options is a direct extension of the option framework from a tabular representation to a linear representation. Linear options are defined over states which are represented as n dimensional feature vector. A linear option is a tuple $(\mathcal{I}, \mu, \beta)$, where $\mathcal{I} \in \mathbb{R}^n$ is the initiation set, $\mu : \mathbb{R}^n \mapsto A$ is the option policy, and $\beta : \mathbb{R}^n \mapsto [0, 1]$ is the termination condition. Sorg and Singh (2010) proposed the idea of linear options as linear representations of options. They proposed an extension of the MDP-LSTD method for policy evaluation in the primitive action case to include options by defining SMDP-LSTD method for behavioral policy evaluation. The solution given by SMDP-LSTD is

$$0 = \sum_k^t \phi_k (r_k + \gamma_k \phi_k'^T \theta_k - \phi_k^T \theta_k) \tag{4}$$

$$\theta^\pi = \left[\sum_{k=1}^t \phi_k (\phi_k - \gamma_k \phi'_k)^T\right]^{-1} \left[\sum_{k=1}^t \phi_k r_k\right] \tag{5}$$

A linear option expectation model of a behavior policy's dynamics is denoted by (F_π, b_π) which attempt to satisfy, for all time steps t.

$$F_\pi \phi_t \approx E\left[\gamma_t \phi'_t | \phi_t\right] \tag{6}$$

$$b_\pi \phi_t \approx E\left[r_t | \phi_t\right] \tag{7}$$

where $F_\pi \phi$ can be interpreted as the expected discounted termination feature vector and $b_\pi \phi$ can be interpreted as the expected discounted reward until termination. A learning and planning agent will need to estimate F_π and b_π from data. Given an input feature vector ϕ, the following recursion defines a LOEM policy evaluation update to the value function parameters θ_k:

$$\theta_{k+1} = \theta_k + \alpha_k (b_\pi^T \phi + \theta_k^T F_\pi \phi - \theta_k^T \phi)\phi \tag{8}$$

LOEM policy evaluation updates is shown to converge to the same behavior value function parameters as does the SMDP-LSTD algorithm.

3 Gradient Descent TD algorithm for linear options

The LSTD algorithms developed can be guaranteed stable under general conditions but they are quadratic in computational complexity. We propose here a gradient descent method (SMDP-TDC) for planning in Reinforcement learning using off-policy TD framework. SMDP-TDC is the direct extension of TDC algorithm (Equations 1, 2 and 3) to the SMDP setting. The option value function is approximated as $V^\pi(s) = \max_o Q^\pi(s, o) \approx \theta^T \phi_s$. We aim to minimize the mean-square projected bellman error ($MSPBE_\theta = \|V_\theta - \Pi T V_\theta\|_D^2$. For the linear option architecture $V_\theta = \Phi\theta$ where Φ is the matrix whose rows are ϕ_s. Π is the projection operator, $\Pi = \Phi(\Phi^T D\phi)^{-1}\Phi^T D$ (where D is the diagonal matrix containing stationary distribution), and T is the SMDP Bellman Operator.

$$\begin{aligned}
&MSPBE_\theta \\
&= \|V_\theta - \Pi T V_\theta\|_D^2 \\
&= \|\Pi(V_\theta - TV_\theta)\|_D^2 \\
&= (\Pi(V_\theta - TV_\theta))^T D(\Pi(V_\theta - TV_\theta)) \\
&= ((V_\theta - TV_\theta))^T \Pi^T D\Pi((V_\theta - TV_\theta)) \\
&= ((V_\theta - TV_\theta))^T D^T \Phi(\Phi^T D\phi)^{-1}\Phi^T D((V_\theta - TV_\theta)) \\
&= (\Phi^T D(TV_\theta - V_\theta))^T (\Phi^T D\Phi)^{-1}(\Phi^T D(TV_\theta - V_\theta)) \\
&= E[\delta\phi]^T E[\phi\phi^T]^{-1} E[\delta\phi]
\end{aligned}$$

Note that

$$E[\phi\phi^T] = \sum_s d_s \phi_s \phi_s^T = \Phi D \Phi^T$$

$$\begin{aligned} E[\delta\phi] &= \sum_s d_s \phi_s \left(r + \gamma^l \sum_{s'} P^o_{ss'} V_\theta(s') - V_\theta(s) \right) \\ &= \Phi^T D(TV_\theta - V_\theta) \end{aligned}$$

Therefore,

$$\begin{aligned} & -\frac{1}{2}\nabla MSPBE_\theta \\ &= E[(\phi - \gamma^l \phi')\phi^T] E[\phi\phi^T]^{-1} E[\delta\phi] \\ &= (E[\phi] - E[\gamma^l \phi' \phi^T]) E[\phi\phi^T]^{-1} E[\delta\phi] \\ &= E[\delta\phi] - \gamma^l E[\phi'\phi^T] E[\phi\phi^T]^{-1} E[\delta\phi] \\ &= E[\delta\phi] - \gamma^l E[\phi'\phi^T] w \end{aligned}$$

where $w \approx E[\phi\phi^T]^{-1} E[\delta\phi]$. Using this form for the gradient of MSPBE, we can write down the gradient descent expression (parameterized by α and β) to minimize MSPBE as follows:

$$\theta_{k+1} = \theta_k + \alpha_k (\delta_k \phi_k - \gamma^l \phi'_k)(\phi_k^T w_k) \tag{9}$$

where w_k is updated as,

$$w_{k+1} = w_k + \beta_k (\delta_k - \phi_k^T w_k)\phi_k \tag{10}$$

and δ_k is the TD error

$$\delta_k = r_k + \gamma^l \theta_k^T \phi'_k - \theta_k^T \phi_k \tag{11}$$

where l is the length of the option being executed.

Convergence proof of SMDP-TDC can be carried out in a similar way as in Sutton et al. (2009). Hence it can be shown that SMDP-TDC converges to the TD fixed point with probability one.

3.1 Convergence Proof for SMDP-TDC

Theorem 1 *Consider the iterations 9 and 10 of the SMDP-TDC algorithm. Let the step-size sequences α_k and β_k, $k \geq 0$ satisfy in this case α_k, $\beta_k > 0$, for all k, $\sum_{k=0}^{\infty} \alpha_k = \sum_{k=0}^{\infty} \beta_k = \infty$, $\sum_{k=0}^{\infty} \alpha_k^2, \sum_{k=0}^{\infty} \beta_k^2 < \infty$ and that $\frac{\alpha_k}{\beta_k} \to 0$ as $k \to \infty$. Further assume that (ϕ_k, r_k, ϕ'_k) is an i.i.d sequence with uniformly bounded second moments. Let $A = E[\phi_k(\phi_k - \gamma^l \phi'_k)^T]$, $b = E[r_k \phi_k]$, and $C = E[\phi_k \phi_k^T]$. Assume that A and C are non singular. Then the parameter vector θ_k converges with probability one to the TD fixed point.*

Proof The proof of this theorem is based on a two time scale difference in the step-size schedule $\{\alpha_k\}$ and $\{\beta_k\}$; refer Borkar (1997) for a convergence analysis of the general two timescale stochastic approximation recursions. The recursions 9 and 10 correspond to the faster and slower recursions respectively. This is because beyond some integer $N_0 > 0 \ \forall \ k$, the increments in 9 are uniformly larger than those in 10 and hence converge faster. Along faster timescale, i.e, the one corresponding to $\{\beta_k\}$, the associated system ODEs corresponds to

$$\dot{\theta}(t) = 0 \tag{12}$$

$$\dot{w}(t) = E[\delta_t \phi_t | \theta(t)] - Cw(t) \tag{13}$$

The ODE 12 suggests that $\theta(t) = \theta$ from the faster timescale perspective. Indeed, recursion 9 can be rewritten as $\theta_{k+1} = \theta_k + \beta_k \xi(k)$, where from 9, $\xi(k) = \left(\frac{\alpha_k}{\beta_k}(\delta_k \phi_k - \gamma^l \phi'_k \phi_k^T w_k)\right) \rightarrow 0$ almost surely as $k \rightarrow \infty$ because $\frac{\alpha_k}{\beta_k} \rightarrow 0$ as $k \rightarrow \infty$. By the Hirsch lemma (Hirsch (1989)), it follow that $\|\theta_k - \theta\| \rightarrow 0$ almost surely as $k \rightarrow \infty$ for the same θ that depends on the initial condition θ_0 of recursion 9.

Consider now the recursion 10. Let $M_{k+1} = (\delta_k \phi_k - \phi_k \phi_k^T w_k) - E[(\delta_k \phi_k - \phi_k \phi_k^T w_k) | \mathcal{F}(k)]$, where $\mathcal{F}(k) = \sigma(w_m, \theta_m, m \leq k; \phi_s, \phi'_s, r_s, s < k)$, $k \geq 1$ are the sigma fields generated by $w_0, \theta_0, w_{m+1}, \theta_{m+1}, \phi_m, \phi'_m, 0 \leq m < k$. It is easy to verify that $M_{k+1}, k \geq 0$ are integrable random variables that satisfy $E[M_{k+1} | \mathcal{F}(k)] = 0, \ \forall k \geq 0$. Also because $r_k, \ \phi_k$ and ϕ'_k have uniformly bounded second moments, it can be seen that

$$E\left[\|M_{k+1}\|^2 | \mathcal{F}(k)\right] \leq c_1 \left(1 + |||w_k||^2 + ||\theta_k||^2\right),$$

for all $k \geq 0$, for some constant $c_1 > 0$ Now consider the ODE pair 12 and 13. Because $\theta(t) = \theta$ from 12, the ODE 13 can written as

$$\dot{w}(t) = E[\delta_t \phi_t | \theta] - Cw(t) \tag{14}$$

Now consider the function $h(w) = E[\delta \phi | \theta] - Cw$, i.e. the driving vector field of the ODE 14. For 14, $w^* = C^{-1} E[\delta \phi | \theta]$ is the unique globally asymptotically stable equilibrium. Let $h_\infty(.)$ be the funciton defined by $h_\infty(w) = \lim_{r \rightarrow \infty} \frac{h(rw)}{r}$. Then $h_\infty(w) = -Cw$. For the ODE $\dot{w}(t) = h_\infty(w(t)) = -Cw(t)$ the origin is a globally asymptotically stable equilibrium because C is a positive definite matrix. The assumptions are now verified and by Borkar and Meyn (2000) Theorem 2.2 we obtain $||w_k - w^*|| \rightarrow 0$ almost surely as $k \rightarrow \infty$.

Consider now the slower time scale recursion 9. Using above, 10 can be rewritten as

$$\theta_{k+1} = \theta_k + \alpha_k (\delta_k \phi_k - \gamma^l \phi'_k \phi_k^T C^{-1} E[\delta_k \phi_k | \theta_k]) \tag{15}$$

Let $\mathcal{G}(k) = \sigma(\theta_m, m \leq k; \phi_s, \phi'_s, r_s, s < k)$ be the sigma fields generated by the quantities $\theta_0, \theta_{m+1}, \phi_m, \phi'_m, 0 \leq m < k$. Let

$$
\begin{aligned}
Z_{k+1} &= (\delta_k \phi_k - \gamma^l \phi'_k \phi_k C^{-1} E[\delta_k \phi_k | \theta_k]) \\
&- E[(\delta_k \phi_k - \gamma^l \phi'_k \phi_k^T C^{-1} E[\delta_k \phi_k \ \theta_k]) | \mathcal{G}(k)] \\
&= (\delta_k \phi_k - \gamma^l \phi'_k \phi_k^T C^{-1} E[\delta_k \phi_k | \theta_k]) \\
&- E[\delta_k \phi_k | \theta_k] - \gamma^l E[\phi'_k \phi_k^T] C^{-1} E[\delta_k \phi_k | \theta_k]
\end{aligned}
$$

It is easy to see that $Z_k, k \geq 0$ are integrable random variables and $E\left[Z_{k+1} | \mathcal{G}(k)\right] = 0, \forall k \geq 0$. Further,

$$
E\left[\left\|Z_{k+1}\right\|^2 | \mathcal{G}(k)\right] \leq c_2 \left(1 + \left\|\theta_k\right\|^2\right), k \geq 0
$$

for some constant $c_2 > 0$ again because r_k, ϕ_k and ϕ'_k have uniformly bounded second moments.

Consider now the following ODE associated with 9

$$
\dot{\theta}(t) = \left(I - E\left[\gamma^l \phi' \phi^t\right] C^{-1}\right) E[\delta \phi | \theta(t)] \tag{16}
$$

Let $\bar{h}(\theta(t))$ be the driving vector field of the ODE 16. Note that

$$
\begin{aligned}
\bar{h}(\theta(t)) &= \left(I - E\left[\gamma^l \phi' \phi^T\right] C^{-1}\right) E\left[\delta \phi | \theta(t)\right] \\
&= \left(C - E\left[\gamma^l \phi' \phi^T\right]\right) C^{-1} E\left[\delta \phi | \theta(t)\right] \\
&= \left(E\left[\phi \phi^T\right] - E\left[\gamma^l \phi' \phi^T\right]\right) C^{-1} E\left[\delta \phi | \theta(t)\right] \\
&= A^T C^{-1} (-A\theta(t) + b)
\end{aligned}
$$

because $E[\delta \phi | \theta(t)] = -A\theta(t) + b$.

Now $\theta^* = A^{-1} b$ can be seen to be the unique globally asymptotically stable equilibrium for 16. Let $\bar{h}_\infty(\theta) = \lim_{r \to \infty} \frac{\bar{h}(r\theta)}{r}$. Then $\bar{h}_\infty(\theta) = -A^T C^{-1} A\theta$. Consider now the ODE

$$
\dot{\theta}(t) = -A^T C^{-1} A \theta(t) \tag{17}
$$

Because C^{-1} is positive definite and A has full rank (as it is nonsingular by assumption), the matrix $A^T C^{-1} A$ is also positive definite. The ODE 17 has the origin as its unique globally asymptotically stable equilibrium. The assumptions are once again verified and the claim follows.

4 Multiple Time Scales and Real Time RL

Autonomous systems make decisions by choosing actions at multiple time scales. Although planning for acting at multiple time scales requires more computational time. Real time systems do not have much time for deliberation. Hence a lot of research has been done to construct just the right number of options by optimizing the computational time around models to accommodate multiple time scales (Stolle and Precup (2002), Wolfe and Barto (2005)). Is it necessary

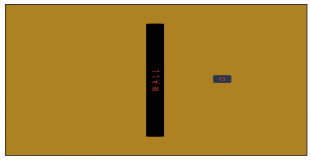

Figure 1: Grid World Domain

to be so careful, or can we do around with many option models? The ideas central to planning is searching and learning while searching. We demonstrate empirically that it is much better to search with many option models, rather than using few handcrafted ones.

We use the sample based search to search through the state space. The search is done by picking up a quasi-random model and sampling the next state, up to a given depth. At the given depth a guess of the value function is used to greedily pick up the best move. Typically a large number of rollouts is performed, and values lower in the search tree are used to update the values higher up in the search tree (Kocsis and Szepesvri (2006), Veness et al. (2010)). Number of rollouts is a parameter that controls the time for deliberation. The experiment was performed on the gridworld domain shown in Figure 1. We use the *hallway going options* where the option models are precoded. Values for all options are learned over 50 trials, using off-policy random behavior. The following six experimental conditions were used to describe what to do on the next time step *a*) Execute the greedy policy using only primitives *b*) Pick the primitive action suggested by the greedy option policy with the hand-crafted options *c*) Pick the primitive action suggested by the greedy policy with random options *d*) Do 100 rollouts of depth 3 using only primitive actions *e*) Do 100 rollouts of depth 3 using primitive actions and crafted options *f*) Do 100 rollouts of depth 3, sampling among 6 options from the random set

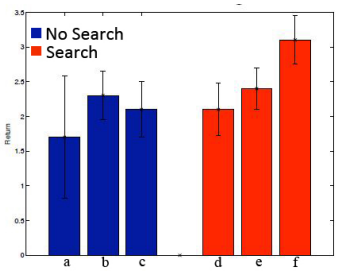

Figure 2: Returns for search

As wee see in Figure 2 search using many option models is better than just using a few handcrafted ones. The results also demonstrate the usefulness of

temporal abstraction and search. Search, and learning using many option model does not effect the time, as the computational complexity of the linear option framework does not depend upon the number of options being used. There are different ways to incorporate search in a real-time system *a*) The planning can be separately threaded *b*) Trying to anticipate what may happen using the model and planning for different states before the sensation is received .

5 Empirical Results

To assess the practical utility of the proposed framework we demonstrate here some experiments where the system learns policy over multiple time scale linear option models updating estimates using the SMDP-TDC algorithm. We also demonstrate the time scales on which the system acts depending upon its position in the environment and its immediate neighbors. The results are demonstrated over 2 domains. The first domain is a grid world domain as shown in Figure 1 and the second domain is the continuous room domain with large dimension feature representation.

5.1 Grid World Domain

The system starts from any state in the state space, it receives a reward of +10 when it reaches the goal state(blue box in Figure 1), while it receives a reward of -3 when it hits the wall and a reward of -1 for all the other transitions. The system has four primitive actions, while acting in the environment with 15% noise. We demonstrate two experimental conditions where option models are composed over multiple time scales and various random policies. The first experiment consists of two time scales - $\beta = 0.5,\ 1$ and biased random policies defined as follows - Pick a probability $x \in (0, 1)$, pick uniformly an action a, the system will pick the action a with probability x and all other actions uniformly with probability $1 - x$. The sytem has same termination conditions, and pick primitive actions with respect to the same policy for all states. Likewise there are 40 policies, hence the system has 80 options available at all states. The second experimental condition consists of 5 termination conditions - $\beta = 0.2, 0.4, 0.50.8, 1$ with the same set of policies and hence it has 200 options available at all states. We plot the average return encounterd by the system averaged over various multiple runs in Figure 3.

Important observations which we note in the experiments are that *a*) Far away from the goal, the best options are the ones biased towards the goal, and with a larger time scale. *b*) Close to the wall, the best options are the ones biased towards the goal and with a shorter time scale. *c*) Still closer, the best options are biased up or down. *d*) The system which does not decide on each time step sometimes hits the wall. *e*) It is much better to search with a lot of random options over multiple time scales rather than a few handcrafted ones .

In Figure 4 we compare the performance of SMDP-TDC method with the primitive TDC method. We also compare the same results without using any

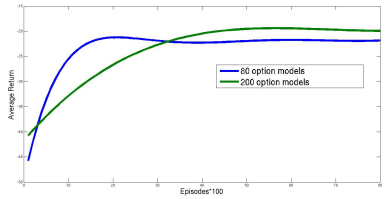

Figure 3: Average Return compared for 2 experimental conditions with options over different multiple time scales (a) 2 timescales of 0.5, 1; and b) 5 time scales of 0.2, 0.4, 0.6, 0.8, 1) by policy learned using SMDP-TDC.

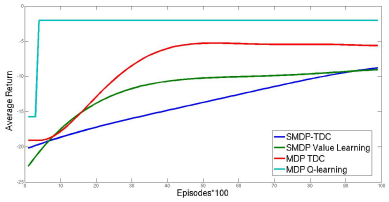

Figure 4: Average Return on the Grid World Domain.

function approximation while following a) SMDP Value Learning, and b)MDP Q learning. We observe the convergence of SMDP-TDC algorithm to a recursive optimal solution. In Figure 4 we observe that option models with multiple time scales reach the goal state more often than its primitive counter part.

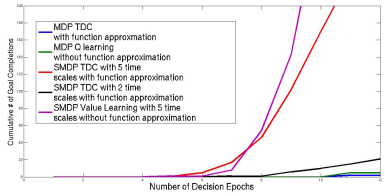
MDP TDC with function approxmation
MDP Q learning without function approximation
SMDP TDC with 5 time scales with function approximation
SMDP TDC with 2 time scales with function approximation
SMDP Value Learning with 5 time scales without function approximation
Cumulative # of Goal Completions
Number of Decision Epochs

Figure 5: Performance in the Grid World Domain

5.2 Continuous Domain

We ran our second set of experiments on a continuous navigation domain as described in Sorg and Singh (2010). The domain (Figure 6), which is 10 $\times$ 10 continuous room world consists of rooms which are separated by walls as shown. In addition the floor of each room is colored with one of the 4 colors (P)urple, (G)reen, (Y)ellow, (B)lue. In figure the floor colors are indicated by the respective first letter, and rooms are seperated by using dashed lines. While

the walls are shown using solid lines.

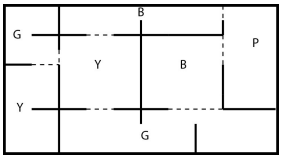

Figure 6: Continuous Room Domain

The agent controls a circular wheel robot for navigation that is capable of observing its current global position $[x, y] \in [0, 10]^2$ and its current global orientation in terms of the an angle $\psi \in [0, 360]$. In addition it can also detect the color of the floor beneath it. It has 3 available primitive actions:`forward` which moves the robot one unit in the forward direction with respect to its current position and orientation, `left` turns the robot towards left by an angle of 30^o at its current location, and `right` which turns the robot towards right with an angle of 30^o at its current location. If the robot hits the wall, its motion is halted in the direction perpendicular to the wall plane. The agent is given a reward of 1 for reaching the purple room on the right side of the domain. After receiving its reward, the agent is transported to the yellow room on the left corner of the building. At all other times, the agent receives a small negative reward of 0.01.

The agent represent states using feature vectors ϕ of 1204 features. The first four features encode the color of the floor in a binary format. The rest of the features are calculated by placing radial basis functions every one step in the $x-y$ direction and every 30^o for ψ. The feature values are calculated using $\phi_i = bexp(-\frac{1}{2}(s-u_i)^T C(s-u_i))$ where $s = [x, y, \psi]$, $b = 10$, $C = diag(\frac{1}{1.2}, \frac{1}{1.2}, \frac{1}{30})$, where u_i are the center of the radial basis functions. To make the feature vector sparse, which can help computationally, any feature that would have had a value $\phi_i < 0.1$, is set to 0.

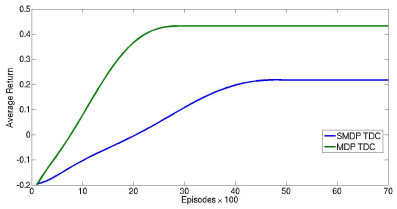

Figure 7: Average Return in the continuous Domain

We tested 2 agents with the above setting, the first agent uses the MDP-TDC learning, while the second agent uses the SMDP-TDC learning using 200 randomly generated option models, with 5 different time scales. The average return obtained is reported in Figure 7. While Figure 8 shows that the number

of goal completions are more when the agent uses option models over multiple time scales as compared to the primitive models.

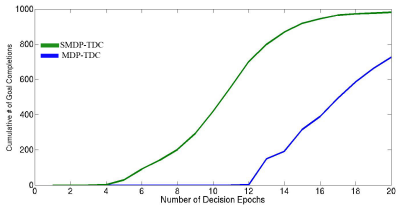

Figure 8: Performance in the continuous Domain

6 Conclusion

Systems which need to act real time, do not have much time for deliberation. For such systems we propose the use of many option models composed over multiple time scale. We also propose an SMDP learning algorithm which is proved to be convergent with the linear option models while achieving a linear time efficiency and memory requirement. We demonstrate our results over the gridworld domain and a continuous room domain. We observe that systems performing while using an approximate representation of the environment states and learning using SMDP TDC algorithm over multiple time scales performs better in all aspects as compared to other methods. They show a recursive optimal behavior which is reasonably good as compared to the optimal solution. Additionally they achieve goals more efficiently as compared to other methods which deliberate at every time step. It will be interesting to see the extension of SMDP TDC algorithm using multiple time-step TD and eligibility traces.

References

Baird, L. (1995). Residual algorithms: Reinforcement learning with function approximation. In *In Proceedings of the Twelfth International Conference on Machine Learning*, pages 30–37. Morgan Kaufmann.

Borkar, V. S. (1997). Stochastic approximation with two time scales. *Systems amp; Control Letters*.

Borkar, V. S. and Meyn, S. P. (2000). The o.d.e. method for convergence of stochastic approximation and reinforcement learning. *SIAM J. CONTROL OPTIM*.

Boyan, J. A. (2002). Technical update: Least-squares temporal difference learning. In *Machine Learning*.

Bradtke, S. J., Barto, A. G., and Kaelbling, P. (1996). Linear least-squares algorithms for temporal difference learning. In *Machine Learning*, pages 22–33.

Hirsch, M. W. (1989). Convergent activation dynamics in continuous time networks. *Neural Netw.*

Kocsis, L. and Szepesvri, C. (2006). Bandit based monte-carlo planning. In *In: ECML-06. Number 4212 in LNCS*.

Sorg, J. and Singh, S. (2010). Linear options. AAMAS '10.

Stolle, M. and Precup, D. (2002). Learning options in reinforcement learning. In *Lecture Notes in Computer Science*.

Sutton, R. and Barto, A. (1998). *Reinforcement Learning: An Introduction*. MIT Press.

Sutton, R. S., Maei, H. R., Precup, D., Bhatnagar, S., Silver, D., Szepesvri, C., and Wiewiora, E. (2009). Fast gradient-descent methods for temporal-difference learning with linear function approximation. In *In Proceedings of the 26th International Conference on Machine Learning*.

Sutton, R. S., Szepesvri, C., and Maei, H. R. (2001). A convergent o(n) algorithm for off-policy temporal-difference learning with linear function approximation. In *Advances in Neural Information Processing Systems 21*. MIT Press.

Veness, J., Ng, K. S., Hutter, M., and Silver, D. (2010). Reinforcement learning via aixi approximation. In *AAAI'10*.

Wolfe, A. P. and Barto, A. G. (2005). Identifying useful subgoals in reinforcement learning by local graph partitioning. In *ICML*.